

Contents

List of Tables	xix
List of Figures	xxi
List of Algorithms	xxv
List of Listings	xxvii
List of Abbreviations	xxix
List of Symbols	xxxiii
1 Introduction	1
1.1 Related Work	4
1.2 Contributions	10
1.3 List of Publications	13
1.4 Structure of this Thesis	16
2 Test Bed	21
2.1 Node-Level Hardware Architectures	21
2.2 Large-Scale Compute Systems	28
2.3 Software Test Bed and Methodology	30
2.4 Test Matrices	32
3 Iterative Sparse Eigenvalue Solvers	35
3.1 Lanczos Method	35
3.2 Kernel Polynomial Method	37
3.3 Chebyshev Filter Diagonalization	39
3.4 Block Jacobi-Davidson QR Method	41
3.5 Summary	43

4	Performance Modeling of Sparse Linear Algebra Building Blocks	47
4.1	The Roofline Performance Model	47
4.2	Application of the Roofline Model to Sparse Linear Algebra	53
4.2.1	(Block-)BLAS-1 Roofline Model	53
4.2.2	Sparse Matrix-Vector Multiplication Roofline Model	53
4.2.3	Sparse Matrix-Multiple Vectors Multiplication Roofline Model	57
4.2.4	General Dense Matrix-Matrix Multiplication Roofline Model	59
4.3	Summary	61
5	A Unified Sparse Matrix Storage Format for High-Performance Sparse Matrix-Vector Multiplication	63
5.1	General Sparse Matrix Storage Formats	63
5.2	The SELL- C - σ Sparse Matrix Storage Format	69
5.3	SELL- C - σ SpMV Implementation	73
5.4	SELL- C - σ SpMV Tuning Factors	78
5.4.1	Chunk Height C	79
5.4.2	Sorting Scope σ	81
5.4.3	OpenMP Scheduling	87
5.5	Unified SELL- C - σ SpMV Performance	91
5.6	Matrix Bandwidth Reduction	101
5.7	Summary	102
6	High-Performance Sparse Matrix-Multiple Vectors Multiplication	105
6.1	Block Vector Storage and SpMMV Implementation	105
6.2	Performance Influence of the Sparse Matrix Storage Format	110
6.3	Performance Influence of Matrix Bandwidth Reduction . .	111
6.4	Summary	112
7	High-Performance Tall & Skinny Matrix-Matrix Multiplication Operations	115
7.1	Tall & Skinny Matrix-Matrix Multiplication Performance .	115
7.2	Tall & Skinny Matrix Transposed-Tall & Skinny Matrix Multiplication	116
7.3	Tall & Skinny Matrix-Small Matrix Multiplication	119
7.4	In-Place Tall & Skinny Matrix-Small Matrix Multiplication .	122
7.5	Summary	123

8	Custom Compute Kernels	125
8.1	KPM Operator	126
8.1.1	Implementation	129
8.1.2	Performance Modeling	130
8.2	ChebFD Operator	135
8.3	Summary	137
9	GHOST: General, Hybrid, and Optimized Sparse Toolkit	139
9.1	Data-Parallel Heterogeneous Execution	139
9.2	Affinity-Aware Tasking	145
9.3	Data Structures	148
9.3.1	Sparse Matrices	148
9.3.2	Dense Matrices	152
9.4	Parallel SpMV	153
9.5	Code Generation	157
9.6	An Example Application Using GHOST	159
9.7	Integration with Other Software	163
9.8	Summary	165
10	Node-Level and Large-Scale Application Performance	167
10.1	KPM Performance Analysis	167
10.2	ChebFD+KPM Performance Analysis	171
10.3	BJDQR Performance Analysis	178
10.4	KPM Energy Analysis	181
10.5	Summary	184
11	Summary	185
12	Outlook	189
	Bibliography	191